

REMARKS

This application contains claims 1-69. New claims 67-69 are hereby added. No new matter has been introduced. Reconsideration is respectfully requested.

Claims 1-9, 17-21, 23-31, 39-43, 45-53 and 61-65 were rejected under 35 U.S.C. 102(e) over Cramer et al. (U.S. Patent 5,946,685). Applicant respectfully traverses this rejection.

Cramer describes a global mount mechanism used in a distributed computing system, which uses a "distributed file system" (col. 3, lines 34-35). A key characteristic of distributed file systems is that, with respect to any file system resource, "one node can act as the server computer for the resource, and the other nodes as client computers" (col. 3, lines 26-28). This client/server model is basic to the mount mechanism that is implemented in Cramer's system, as can be seen in all of the figures that illustrate its operation (such as Figs. 3A, 3B, 7, 8A-8D, 10 and 12). Although the classification of client and server can vary over time, there is always one, clearly-defined server for any given file system resource at any given time.

The present patent application distinguishes clearly between distributed file systems, such as that described by Cramer, and parallel file systems. Distributed file systems are defined in the specification as follows (page 3, line 26 - page 4, line 12):

...Distributed file systems... allow a user on a client computer connected to a network to access and modify data stored in files on a file server. When a user accesses data on the file server, a copy of the

data is stored, or cached, on the client computer, and the user can then read and modify the copy... All of these distributed file systems, however, are still essentially single-node systems, in which a particular server controls any given file. The DMAPI and data management applications for such distributed file systems are essentially server functions and are not distributed among the client nodes.

A parallel file system, on the other hand "allows applications on multiple nodes to share file data, without mediation of a file server as in distributed file systems" (page 4, lines 17-19). The parallel file system enables all nodes in a cluster to access the same file data concurrently (page 15, lines 12-21).

Mounting and unmounting of file systems in a parallel file system environment creates problems for data management applications that do not exist in distributed file system environments. As noted in the specification (page 7, line 19 - page 8, line 4):

It is a characteristic of parallel file systems that multiple instances of the same file system may be mounted on different nodes in the cluster... In single-node systems [including distributed file systems]... the DMAPI reports a single mount event for each file system, as well as a single preunmount and unmount event. Once the file system is unmounted from the node, there can be no further events from this file system until the file system is mounted again. In the parallel file system of the present invention, however, the DMAPI is capable of handling multiple mount events – from all of the nodes on which the file

system is mounted – as well as multiple preunmount and unmount events.

Thus, in distinction to distributed file systems, for example, the parallel file system DMAPI must provide methods to deal with the possibility that “there is not a predictable serial relationship between all the mount, preunmount and unmount events of each file system” (page 27, lines 3-5). The invention recited in the claims of the present patent application provides a solution to these problems.

Claim 1 is drawn to a method for managing data storage, in which a session of a data management (DM) application is initiated on one node in a cluster, and a request is submitted to parallel file system software on a second node to mount a file system on the second node. In response, a mount event message is sent from the second node to the first node for processing by the DM application on the first node.

In rejecting the claims in the present patent application, it appears that the Examiner ignored the material claim limitations relating to the use of a parallel file system. These limitations apply to all of the claims in the present patent application. Cramer makes no mention of any sort of parallel file system, and therefore cannot serve as the basis for a rejection of the claims on the basis of anticipation. As stated in MPEP 2131:

To anticipate a claim, the reference must teach every element of the claim. “The identical invention must be shown in as complete detail as is contained in the ... claim.” *Richardson v. Suzuki Motor Co.*, 868 F.2d 1226, 1236, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989).

As noted earlier, the distinction between Cramer's distributed file system and the parallel file system recited in the claims of the present patent application is substantive, and must be taken into account in examining the claims. Furthermore, there is no teaching or suggestion in the prior art that would have led a person of ordinary skill to make the necessary changes to mount techniques that are used in distributed file systems in order to meet the needs of data management applications in a parallel file system, as recited in claim 1. Therefore, claim 1 is believed to be patentable over Cramer.

Even if it were to be conceded that claim 1 is anticipated by Cramer, there is still insufficient basis in Cramer to support the novelty rejection of the dependent claims. For example, claims 8 and 9, which depend from claim 1, recite the use of novel flags in mount and unmount event messages in order to permit the first node to determine the source of mount and unmount requests. These flags are needed in order to deal with ambiguities that may arise in a parallel file system, in which multiple mount, preunmount and unmount events may occur in an unpredictable order (page 27, lines 3-31). Distributed file systems have no such need.

Cramer, in fact, neither teaches nor suggests the use of flags in event messages of any sort. Moreover, the passage in Cramer that was alleged by the Examiner to anticipate claim 8 (col. 12, lines 22-31) has nothing to do with "determining... whether the further request [to unmount a file system instance] was submitted on the first node or on another one of the nodes," as recited in the claim. Similarly, the passage cited against claim 9 (col. 9, lines 47-53) relates to determining the server that is to fulfill a given mount request, and has nothing to do

with "determining... whether the request was submitted on the first node or on another one of the nodes." Comparable arguments may be made against the rejection of the other dependent claims. For the sake of brevity, however, and in view of the patentability of claim 1, these arguments are omitted here.

Thus, to summarize, Applicant respectfully submits that claims 1-9 and 17-21 are patentable over Cramer.

Claims 23 and 45 respectively recite computing apparatus and a computer software product, which operate on principles similar to the method of claim 1, in the explicit context of parallel file systems. Therefore, for the reasons stated above, claims 18 and 35 are likewise believed to be patentable over Cramer. In view of the patentability of these independent claims, dependent claims 24-31, 39-43, 46-53 and 61-65 are believed to be patentable, as well.

Claims 10-16, 32-38 and 54-60 were rejected under 35 U.S.C. 103(a) over Cramer in view of Dugan et al. (U.S. Patent 6,363,411), while claims 22, 44 and 66 were rejected under 35 U.S.C. 103(a) over Cramer in view of Dugan and further in view of Vahalia et al. (U.S. Patent 6,192,408). In view of the patentability of independent claims 1, 18 and 35, from which these claims depend, Applicant believes these dependent claims to be patentable, as well, over the cited art.

Applicant has added new claims 67-69 in order to more clearly and completely recite aspects of the present invention. These new claims depend from claims 1, 23 and 45, respectively. They state that a request to mount one of the file systems is submitted by a user application running on a second node in the cluster. In response to this request, a mount event message is sent from the second node to a first node, on which a data management application session is running. This is a

typical paradigm in a cluster that uses a parallel file system according to the present invention. It is shown, for example, in Figs. 2 and 3 of the present patent application, and is described on page 19, lines 27-29, page 21, lines 1-4, and page 24, line 19 - page 25, line 17, of the specification.

In rejecting claims 1, 23 and 45, the Examiner associated the step of "initiating a session of a data management application" with the function of issuing a global mount command at a requesting node, described by Cramer in col. 9, lines 22-23. In other words, according to the Examiner's interpretation, the "first node" of claims 1, 23 and 45 is identified with Cramer's "requesting node," which requests that a file system resource be mounted. The Examiner then associated "receiving a request... at a second one of the nodes" with the function of requesting that the server for the mount point in question generate a file object (col. 9, lines 47-53, in Cramer), so that the "second node" recited in claims 1, 23 and 45 is identified with Cramer's server node. In Cramer's system, the server then mounts the file system and sends information to the requesting node (col. 9, lines 59-60).

In new claims 67-69, on the other hand, the first node has nothing to do with submitting or fulfilling the mount request. Rather, the mount request is submitted by a user application running on the second node, and the second node then mounts the requested file system. The second node sends an event message to notify the first node of the mount request by virtue of the data management application session running on the first node, not because the first node has requested the mount. Cramer neither teaches nor suggests this sort of model, nor would there be any sort of motivation to implement this sort of model in a distributed

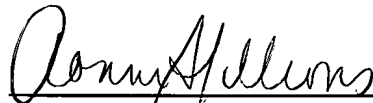
file system of the sort that Cramer describes. Thus, claims 67-69 are believed to be patentable over the cited art.

Applicant has studied the additional references that were made of reference by the Examiner, and believes the claims in the present patent application to be patentable over these references, as well, whether the references are taken individually or in any combination.

Applicant believes the amendments and remarks presented hereinabove to be fully responsive to all of the grounds of rejection raised by the Examiner. In view of these amendments and remarks, applicant respectfully submits that all of the claims in the present application are in order for allowance. Notice to this effect is hereby requested.

Respectfully submitted,

BROWDY AND NEIMARK, P.L.L.C.
Attorneys for Applicant

By 
Ronni S. Jillions
Registration No. 31,979

RSJ:ft
Telephone No.: (202) 628-5197
Facsimile No.: (202) 737-3528
G:\BN\C\colb\Loy3\PTO\19March04AMD.doc